arXiv:1407.3957v2 [cs.GT] 17 Sep 2014

Efficiency of Truthful and Symmetric Mechanisms in One-sided Matching*

Marek Adamczyk[1], Piotr Sankowski[2], and Qiang Zhang[2]

[1]Sapienza University of Rome, Italy, adamczyk@dis.uniroma1.it
[2]Institute of Informatics, University of Warsaw, sank,qzhang@mimuw.edu.pl

September 18, 2014

Abstract

We study the efficiency (in terms of social welfare) of truthful and symmetric mechanisms in one-sided matching problems with *dichotomous preferences* and *normalized von Neumann-Morgenstern preferences*. We are particularly interested in the well-known *Random Serial Dictatorship* mechanism. For dichotomous preferences, we first show that truthful, symmetric and optimal mechanisms exist if intractable mechanisms are allowed. We then provide a connection to online bipartite matching. Using this connection, it is possible to design truthful, symmetric and tractable mechanisms that extract 0.69 of the maximum social welfare, which works under assumption that agents are not adversarial. Without this assumption, we show that Random Serial Dictatorship always returns an assignment in which the expected social welfare is at least a third of the maximum social welfare. For normalized von Neumann-Morgenstern preferences, we show that Random Serial Dictatorship always returns an assignment in which the expected social welfare is at least $\frac{1}{e}\frac{\nu(\mathcal{O})^2}{n}$, where $\nu(\mathcal{O})$ is the maximum social welfare and n is the number of both agents and items. On the hardness side, we show that no truthful mechanism can achieve a social welfare better than $\frac{\nu(\mathcal{O})^2}{n}$.

1 Introduction

We study the efficiency of mechanisms in one-sided matching problems, where the goal is to allocate n *indivisible* items to n unit-demand *rational* agents having *private* preferences over items. Agents are rational, i.e., they would like to be assigned to the best items according to their private preferences. The problem essentially captures variants of practical applications such as allocating houses to residents, assigning professors to courses and so on. In this paper, we mainly focus on *cardinal preferences* in which agents have values for different items. A practical setting would be that residents have values for different houses. A *mechanism* maps preferences that agents report to a matching, which is a one-to-one mapping between agents and items. Throughout the paper, depending on the context, we use sometimes term *matching* and sometimes *assignment*, but they always mean essentially the same. One immediate question arises: if there exist mechanisms in

*This work was partially supported by ERC StG project PAAl 259515, FET IP project MULTIPEX 317532, and NCN grant N N206 567940.

which no agent could benefit by deviating from reporting his true preference regardless the preferences reported by other agents? Such mechanisms are often called *truthful* mechanisms. The question was answered in [15], where it was shown that there exists only one *truthful, nonbossy* and *neutral* mechanism. A mechanism is nonbossy if an individual agent cannot change the output of the mechanism without changing his assignment. A mechanism is neutral if the mechanism is independent of the identities of items, e.g., the assignment get permuted accordingly when the items are permuted. The unique mechanism works as follows. First, agents are sorted in a fixed order, and then the first agent chooses his favorite item, the next agent chooses his favorite item among remaining items, etc. When the fixed order is picked uniformly among all possible orderings, the resulted mechanism is called *Random Serial Dictatorship (RSD)*.

Besides the truthfulness, an important issue left is to understand the efficiency of mechanisms in one-sided matching problems. The efficiency of a mechanism is defined as the social welfare of the assignment the mechanism returns. Zhou [17] confirmed Gale's conjecture by showing that there is no *symmetric, Pareto optimal* and *truthful* mechanism for general preferences. A mechanism is symmetric if agents are treated equally if they report the same preferences. A mechanism is Pareto optimal if the mechanism never outputs an assignment that the social welfare could be improved without hurting any agent. It is well-known that RSD is truthful and ex post efficient, i.e., it never outputs Pareto dominated outcomes.

We observe that there is few work that study the efficiency of RSD. The main reason is that its average social welfare could be even $O(n)$ far away from the optimal social welfare if the preferences of agents for items are unrestricted. It happens when assigning a particular item to a particular agent contributes most of the optimal social welfare. However, in RSD it is possible that the agent only gets that item with a probability of $1/n$. In this paper, we circumvent this problem by considering smaller but still rich domains of preferences. The first type of preferences we consider is *dichotomous preferences*, where agents have binary preferences over items. We shall call this setting simply *dichotomous*. Dichotomous preferences are fairly natural in assignment problems. For example, professors indicate the courses they like or dislike to teach, or workers choose the working shifts they want. The goal here is to design good mechanisms to assign courses/shifts to professors/workers. One can model these problems with bipartite graphs: workers on one side, shifts on the other, an edge indicates whether a worker wants to participate in a particular shift. Then one can find a maximum matching in the graph to optimize the total value of the assignment. It is shown in [5] that with some careful tie-breaking rule, finding a maximum matching yields a truthful mechanism. However, such mechanisms fail to capture the symmetry. To make this approach symmetric, one could find all maximum matchings and randomly choose one. Note that it implies that Zhou's impossibility result does not pertains to dichotomous preferences. However, since finding all maximum matchings in bipartite graphs is $\#P$-complete, we conjecture that it is computationally infeasible to design truthful and symmetric mechanisms that obtain optimal welfare. Therefore, we turn our attention to investigate how well mechanisms can approximate the maximum social welfare. By the connection to the online bipartite matching problem [11, 12], we get the following result:

Result 1. *In dichotomous setting there exists a truthful and symmetric mechanism that is a* 0.69*-approximation to the maximum social welfare.*

Due to the limitation of such mechanisms, next we show that RSD also obtains a constant approximation for dichotomous preferences.

Result 2. *Random Serial Dictatorship in dichotomous setting returns an assignment in which the expected social welfare is a 3-approximation of the maximum social welfare.*

The second type of preferences we consider is *normalized von Neumann-Morgenstern preferences*, where the value of agent i for item j lies in $[0, 1]$. We shall call this setting simply *normalized*. In this setting our result gives asymptotically tight description of the social welfare achieved by RSD.

Result 3. *In normalized setting with n agents and n items, Random Serial Dictatorship returns a matching which expected social welfare is at least $\frac{1}{e}\frac{\nu(\mathcal{O})^2}{n}$, where $\nu(\mathcal{O})$ is the maximum social welfare.*

This result implies that RSD achieves an $\sqrt{e \cdot n}$-approximation of the optimal social welfare in unit-range preferences, i.e., when $\max_i v_a(i) = 1$, $\min_i v_a(i) = 0$. Recently [6] presented an $O(\sqrt{n})$-approximation for RSD in unit-range setting.

Finally, we complement the above result with the following upper-bound.

Result 4. *Given n, for any $k = 1, \ldots, n$ and for any $\epsilon > 0$ there exist an instance of one-sided matching problem with normalized von Neumann-Morgenstern preferences where $\nu(\mathcal{O}) = k$ and no truthful mechanism can achieve expected social welfare better than $\frac{k^2}{n} + \epsilon$, where k is the optimal social welfare.*

1.1 Related work

Here we only mention the most relevant work on one-sided matching problems. For more details, we refer the reader to surveys [13, 14]. One-sided matching problems modeled in [9] gave a market-like procedure to produce efficient assignments. There, the procedure is Pareto optimal but not truthful. Gale and Shapley [7] considered a similar problem, the marriage problem, but they turned attention to the incentive issues on whether agents would or would not reveal their private preferences. In [8] authors were asking about existence of good mechanisms when preferences are also considered. Zhou [17] answered this question by showing that there is no *symmetric, Pareto optimal* and *truthful* mechanism. Between ex-ante Pareto optimality and ex-post Pareto optimality, Bogomolnaia and Moulin [2] introduced a new concept called *ordinal efficiency*. They gave a *probabilistic serial* mechanism that always returns ordinal efficient assignments. However, the probabilistic serial mechanism is not truthful. Bhalgat et al. [1] studied the efficiency of RSD in a more restricted setting than ours, where agents have values of $\frac{n-j+1}{n}$ for their jth favorite item. Chakrabarty and Swamy [4] introduced the notion of rank approximation to measure the social welfare under ordinal preferences. One-sided matching problems with dichotomous preferences were studied by Bogomolnaia and Moulin [3]. They used the Gallai-Edmonds decomposition of bipartite graphs to characterize the (most) efficient assignments.The most related work to ours is that Filos-Ratsikas et al. [6] independently gave the similar approximation ratio of RSD under unit-range preferences while our results applies to more general settings.

Cardinal preferences enable agents to explicitly express how much they prefer each item, while this can not be done in ordinal preferences. The space of cardinal preferences could be shown to be the same as the space of von Neumann-Morgenstern preferences. In addition, the normalization of preferences is a standard procedure, see [10]. Besides the literature of operational research and decision theory, normalized von Neumann-Morgenstern preferences are widely used to model individual behavior in game-theoretical settings.

2 Preliminaries

The model We model one-sided matching problems as bipartite graphs. In a bipartite graph, its left side is a set A of agents and its right side are a set I of indivisible items. We assume $|A| = |I| = n$ and each agent is matched to exactly one item. For each agent $a \in A$ and each item $i \in I$, there is an edge (a, i) representing a possible allocation of item i to agent a. The preference of agent a for item i is denoted by $v_a(i)$, which is the value that agent a has for item i. We consider two different types of preferences, *dichotomous preferences* and *normalized von Neumann-Morgenstern preferences*. In dichotomous preferences, it holds that $v_a(i) \in \{0, 1\}$, while in normalized von Neumann-Morgenstern preferences, it holds that $v_a(i) \in [0, 1]$. In dichotomous case we shall say shortly that agent a *1-values* item i, if $v_a(i) = 1$, instead of clunky "agent a has value 1 for item i"; the same with value 0.

We say $v_a(\cdot)$ is the preference profile of agent a. Denote by $\mathcal{V}$ the set of all possible preference profiles of a single agent: for dichotomous preferences $\mathcal{V} = \{0, 1\}^I$, for normalized von Neumann-Morgenstern preferences $\mathcal{V} = [0, 1]^I$. Preference profiles of all agents are denoted by $v_A = (v_a)_{a \in A} \in \mathcal{V}^A$; by $v_{-a} = (v_{a'})_{a' \in A \setminus a}$ we denote all profiles except of agent a's. By (v'_a, v_{-a}) we denote agents' preferences with a's preference changed from v_a to v'_a; if (v'_a, v_{-a}) is an argument of a function, then we skip writing double brackets. Consider a set of items $I' \subseteq I$ and suppose that agent a values items $i_1, \ldots, i_k \in I'$ equally and more than any other item in I'. We say that items $i_1, \ldots, i_k$ are *favorite* items of agent a in I'.

We call matrix $p_A = (p_a)_{a \in A}$, where $p_a = \left(p_a^i\right)_{i \in I}$, a feasible *matching* if the following conditions hold: 1) for any $a \in A$ and $i \in I$, $p_a^i \in \{0, 1\}$; 2) for any $a \in A$, $\sum_{i \in I} p_a^i = 1$; 3) for any $i \in I$, $\sum_{a \in A} p_a^i = 1$. Given a feasible matching p_A, we say item i is matched to agent a if $p_a^i = 1$. Thus, the value of agent a for the matching p_A is given by $v_a \cdot p_a = \sum_{i \in I} v_a(i) p_a^i$, where $\cdot$ is an operator of the vector product. The social welfare of the matching p_A is given by $\nu(p_A) = \sum_{a \in A} v_a \cdot p_a$.

From each agent $a \in A$ mechanism $\mathcal{M}$ collects declarations $d_a \in \mathcal{V}$ about his preference profile — we overload notations here a bit, since vector d_a does not always have to be declared completely, i.e., when some of the items are already matched, then the mechanism does not ask a about values for these items. Of course, the connection between true valuations $v_a \in \mathcal{V}$ and declarations $d_a \in \mathcal{V}$, which $\mathcal{M}$ collects, depends heavily on the mechanism $\mathcal{M}$ itself. Mechanism $\mathcal{M}$ maps agents declarations d_A to a feasible matching $\mathcal{M}_A(d_A)$ (i.e., the p_A matrix); $\mathcal{M}_a(v_A)$ denotes the allocation to agent a (i.e., the p_a vector). Mechanism $\mathcal{M}$ might be randomized, and then matching $\mathcal{M}_A(d_A)$ is a random matrix, and allocation $\mathcal{M}_a(v_A)$ is a random vector as well. In this case, $\mathbb{E}[\nu(\mathcal{M}_A(d_A))]$ is the expected social welfare of mechanism $\mathcal{M}_A$, but since all of the mechanisms we analyze are randomized, we shall call it just social welfare.

We measure the performance of the mechanism by comparing the social welfare it produces with the optimal social welfare $\nu(\mathcal{O}(v_A))$, where $\mathcal{O}(v_A)$ denotes a matching that maximizes the social welfare when preferences are given by v_A. Note that $\mathcal{O}(v_A)$ can be seen as a maximum weight matching in the graph $G = (A \cup I, A \times I)$ where weight of edge (a, i) is equal to $v_a(i)$. For simplicity however, throughout the paper we shall just write $\mathcal{O}$, instead of $\mathcal{O}(v_A)$.

A mechanism $\mathcal{M}$ is *truthful*, if for every $a \in A$, every $v_A \in \mathcal{V}^A$ and every $v'_a \in \mathcal{V}$, it holds that (even when the mechanism is randomized)

$$v_a \cdot \mathcal{M}_a(v_A) \geq v_a \cdot \mathcal{M}_a\left(v'_a, v_{-a}\right).$$

A mechanism $\mathcal{M}$ is *symmetric* if for every $a_1, a_2 \in A$, every $d_A \in \mathcal{V}^A$ such that $d_{a_1} = d_{a_2}$, it holds

that $\mathbb{E}\left[\mathcal{M}_{a_1}\left(d_A\right)\right] = \mathbb{E}\left[\mathcal{M}_{a_2}\left(d_A\right)\right]$, i.e., agents with identical declarations have the same (expected) value for the allocation.

RSD and iterative analysis Now let us give the formal description of the Random Serial Dictatorship (RSD) mechanism. RSD first picks an ordering of agents uniformly at random and then asks agents to choose sequentially with respect to the order. We assume that agents are rational, i.e., they will always choose the best items among the unmatched items. Ties are broken randomly, i.e., when agent a is asked in RSD and his favorite items are i_1 and i_2 among unmatched items, agent a will chose items i_1 and i_2 with an equal probability. This is an important assumption for the analysis of RSD with dichotomous preferences. If we would like to analyze RSD when agent would always deterministically choose among the best items, then the competitive ratio guarantees and lower bounds from von Neumann-Morgenstern preferences would apply.

Let us observe a property of RSD that is important for our analysis. Instead of thinking that a random ordering is fixed before any agent is considered sequentially, we can think that RSD chooses an agent randomly from remaining agents in each step. It is easy to see that agents are considered in the same random order in both cases.

RSD is iterative in nature, and so is the analysis. Let us index its time-steps by t, which ranges from 0 to n. $t = 0$ indicates the moment after sorting the agents, but before asking first agent to choose. Let $\mathcal{R}^t$ represent the (partial) matching constructed by RSD after first t steps. Then $\nu\left(\mathcal{R}^t\right)$ represents the social welfare obtained after first t steps; in particular $\nu\left(\mathcal{R}^0\right) = 0$. As RSD is being executed, the set of unmatched agents and the set of available items are gradually decreasing. Let A^t and I^t be the set of unmatched agents and the set of available items after step t. For example, $A^0 = A$ and $I^0 = I$. As the sets A^t and I^t are being modified, we also keep track of the way in which $\nu\left(\mathcal{O}\right)$ is being changed (recall that $\mathcal{O}$ denotes a matching that maximizes the welfare). More precisely, we start with $\nu\left(\mathcal{O}^0\right) = \nu\left(\mathcal{O}\right)$. Suppose that at step t, RSD asks agent a to choose and then a picks item i, then $\nu\left(\mathcal{R}^t\right) = \nu\left(\mathcal{R}^{t-1}\right) + v_a\left(i\right)$. We remove a from A^{t-1} and i from I^{t-1}, e.g., $A^t = A^{t-1} - \{a\}$ and $I^t = I^{t-1} - \{i\}$. In addition, we also remove welfare contributed by a and i from $\nu\left(\mathcal{O}^{t-1}\right)$. Certainly, when $t = n$, then $\nu\left(\mathcal{O}^n\right) = 0$, while $\nu\left(\mathcal{R}^n\right)$ is the social welfare obtained by RSD.

Sequence $\left\{\nu\left(\mathcal{R}^t\right)\right\}_{t\geq 0}$, which represents the increasing welfare of RSD, is a random process. Moreover, $\mathbb{E}\left[\nu\left(\mathcal{R}^n\right)\right]$ represents the expected social welfare returned by RSD. The sequence $\left\{\nu\left(\mathcal{O}^t\right)\right\}_{t\geq 0}$, which represents how the optimal social welfare is affected by the random choices within RSD, is a random process as well. Therefore, we want to describe a relation between $\mathbb{E}\left[\nu\left(\mathcal{R}^n\right)\right]$ and $\nu\left(\mathcal{O}^0\right)$, and to do so we deploy theory of martingales.

Martingales Below we only introduce notions and properties that we use later in the paper. For a thorough treatment of martingale theory see [16].

Definition 1. Consider a random process $\left(X^t\right)_{t=0}^n$. Suppose we observe first k steps of the process, and let $\mathcal{H}^k$ denote the information we have acquired in steps $0, 1, \ldots, k$. Expected value of X^{k+1}, conditioned on the information we have from steps 0 to k, is formally presented as $\mathbb{E}\left[X^{k+1}\middle|\,\mathcal{H}^k\right]$. If for any $k = 0, \ldots, n-1$, we have $\mathbb{E}\left[X^{k+1}\middle|\,\mathcal{H}^k\right] = X^k$, then the process is called a *martingale.*

In other words, the process does not change on expectation in one step. We shall also consider a *sub-martingale* $\left(X^t\right)_{t=0}^n$ which satisfies $\mathbb{E}\left[X^{k+1}\middle|\,H^k\right] \geq X^k$ instead of equality in the above definition.

Theorem (Doob's Stopping Theorem)**.** *Let* $\left(X^t\right)_{t=0}^n$ *be a martingale, respectively sub-martingale. For any* $k = 0, 1, \ldots, n$ *it holds that* $\mathbb{E}\left[X^k\right] = \mathbb{E}\left[X^0\right]$*, respectively* $\mathbb{E}\left[X^k\right] \geq \mathbb{E}\left[X^0\right]$*.*

The above is not the Doob's theorem in its full generality, but rather the simplest variant that still holds in our setting.

3 Dichotomous preferences and online bipartite matching

In this section, we establish a connection between one-sided matching with dichotomous preferences and online bipartite matching. A similar connection was also presented in [1].

Consider a variant of online bipartite matching. We are given a bipartite graph $G = (A \cup B, E)$, where one side A of the graph is given, while vertices from other side B and edges between A and B are unknown. Suppose that vertices from B arrive one by one, and upon the arrival of vertex $b \in B$, all edges adjacent to b are revealed. On vertices of A there is an ordering σ given by a random permutation. Consider RANKING algorithm that upon arrival of vertex $b \in B$ it matches b to the unmatched neighbor in $a \in A$ with the highest ranking $\sigma(a)$. In their seminal paper, Karp et al. [11] have proven that this algorithm constructs a matching of expected size at least $\left(1 - \frac{1}{e}\right) OPT$, where OPT is the offline optimum, and the bound holds even if the vertices of B arrive in an adversarial order. Furthermore, Mahdian and Yan [12] have shown that the performance of RANKING algorithm is even better when the order of vertices in B is also given by a random permutation:

Theorem . *Given that the vertices in* B *arrive uniformly at random and the order of vertices in* A *is random, RANKING algorithm constructs a matching of expected size at least* $0.69 \cdot OPT$*, where* OPT *is the offline optimum.*

Now let us see consider the following mechanisms for one-sided matching with dichotomous preferences. Given the agents and items, mechanism RSD* generates a random ordering on agents and a random ranking on items. RSD* considers agents one by one according to the random ordering. Suppose that agent a is considered at step τ and let $d_a(\cdot)$ be the preference reported by agent a. Denote by I^τ the set of items yet unmatched at step τ. If agent a 1-values any unmatched item, RSD* assigns agent a an item with the highest rank among all remaining items. Otherwise, RSD* assigns nothing to agent a. Finally, RSD* matches any unmatched items to unmatched agents. Truthfulness of RSD* follows from the observation that τ as well as I^τ are independent of a's declaration d_a. More precisely, the moment τ is given only by a random permutation of agents, while set I^τ depends on the permutation of agents and declarations $d_{a'}$ of agents a' that came before a. Therefore, if a declares $d_a(i) = 1$ for item i such that $v_a(i) = 0$, then he can only increase the probability that at moment τ he is matched to a 0-valued item. Analogically, if a declares $d_a(i) = 0$ for item i such that $v_a(i) = 1$, then he can only decrease the probability that at moment τ he is matched to a 1-valued item. Suppose now that agent a has 0 value for all items in I^τ. In this case agent gains nothing regardless of what his declarations are. An agent that reports truthfully in this case, we call *non-adversarial*. Since RSD* is guided by two random permutations, the symmetry of the mechanism is clear.

Theorem 1. *Assuming that agents are non-adversarial, RSD* is a truthful and symmetric mechanism that achieves* 0.69*-approximation to the maximum social welfare in one-sided matching problems with dichotomous preferences.*

Algorithm 1: RSD*(A, I)

Let random permutation $\sigma : \{1, ..., n\} \mapsto \{1, ..., n\}$ be the ranking of items;
For each agent $a \in A$ in random order:
    ask agent a about his preference profile $d_a \in \mathcal{V} = \{0, 1\}^I$;
    if there is no unmatched item i such that $d_a(i) = 1$, then discard agent a;
    otherwise, assign a to unmatched item i that has the highest rank $\sigma(i)$;
Match any unmatched items to unmatched agents anyhow.

One can imagine that sometimes an agent can be adversarial, and he would not admit that he does not value any of the remaining items. To address this issue, in the next section we present an analysis of RSD mechanism in which every agent can be adversarial.

4 Dichotomous preferences and RSD

Theorem 2. *Random Serial Dictatorship always returns an matching in which the expected welfare is at least* $\frac{1}{3}\nu(\mathcal{O})$ *in one-sided matching problems with dichotomous preferences.*

Proof. Recall, $\mathcal{O}$ is an optimal matching. Let A^t be the set of agents remaining after t steps, let I^t be the set of remaining items, and $\mathcal{O}^t \subset \mathcal{O}$ is what remains from optimal solution after t steps of RSD. Also, $\mathcal{R}^t$ is the partial matching constructed by RSD after t steps, and $\nu\left(\mathcal{R}^t\right)$ be its welfare. For an agent a let $\mathcal{O}_a \in I$ be the item to which a is matched in $\mathcal{O}$.

Let Y^t be the set of agents who are matched to an item in $\mathcal{O}^t$ which they value 1, i.e., $\left\{ a \in A^t \middle| v_a(\mathcal{O}_a) = 1 \right\}$. Therefore $\left|Y^t\right| = \nu\left(\mathcal{O}^t\right)$ for every t. It can happen that at time t, an agent does not 1-value any of remaining items I^t, even though he could have 1-valued some of the items in I^0. Thus let $Z^t \subseteq A^t$ be the agents who 0-value all items in I^t. Let us denote $y^t = \left|Y^t\right|$ and $z^t = \left|Z^t\right|$ for brevity.

Consider step $t+1$ of RSD, and assume we have all information available after first t steps, represented by $\mathcal{H}^t$. Let a be the agent who is to make his choice in this step, and let i be the item a chooses. Agent a does not belong to Z^t with probability $1 - \frac{z^t}{n-t}$, and if this happens, then for sure $v_a(i) = 1$, which adds 1 to the welfare of RSD, i.e., $\nu\left(\mathcal{R}^{t+1}\right) = \nu\left(\mathcal{R}^t\right) + 1$. Hence $\mathbb{E}\left[\nu\left(\mathcal{R}^{t+1}\right) \middle| \mathcal{H}^t\right] = \nu\left(\mathcal{R}^t\right) + 1 - \frac{z^t}{n-t}$.

Now let us analyze the expected decrease $\nu\left(\mathcal{O}^t\right) - \nu\left(\mathcal{O}^{t+1}\right)$. Suppose that agent a does not belong to Z^t, again with probability $1 - \frac{z^t}{n-t}$. Edge (a, i) is adjacent to at most two 1-value edges in $\mathcal{O}^t$, since $\mathcal{O}^t$ is a feasible matching. Thus when $a \notin Z^t$, then $\nu\left(\mathcal{O}^t\right) - \nu\left(\mathcal{O}^{t+1}\right)$ is at most 2 . Now suppose that agent a belongs to Z^t, which happens with probability $\frac{z^t}{n-t}$. Since $v_a(i) = 0$, then a is **not** adjacent to any 1-value edge in $\mathcal{O}^t$, and i may be adjacent to at most one such edge since agent a choose an item randomly from unmatched items. Therefore, when $a \in Z^t$, then $\nu\left(\mathcal{O}^t\right) - \nu\left(\mathcal{O}^{t+1}\right)$ is at most 1. Hence, together with noting that $\frac{z^t}{n-t} + \frac{y^t}{n-t} \leq 1$, we can conclude that the expected decrease $\nu\left(\mathcal{O}^t\right) - \nu\left(\mathcal{O}^{t+1}\right)$ is:

$$\mathbb{E}\left[\nu\left(\mathcal{O}^t\right) - \nu\left(\mathcal{O}^{t+1}\right) \middle| \mathcal{H}^t\right] \leq 2 \cdot \left(1 - \frac{z^t}{n-t}\right) + \frac{z^t}{n-t} \cdot \frac{y^t}{n-t} \leq 3 \cdot \left(1 - \frac{z^t}{n-t}\right).$$

Since $\mathbb{E}\left[\nu\left(\mathcal{R}^{t+1}\right)\middle|\,\mathcal{H}^t\right] = \nu\left(\mathcal{R}^t\right) + 1 - \frac{z^t}{n-t}$, we get that

$$\mathbb{E}\left[\nu\left(\mathcal{O}^t\right) - \nu\left(\mathcal{O}^{t+1}\right)\middle|\,\mathcal{H}^t\right] \leq 3 \cdot \left(1 - \frac{z^t}{n-t}\right) = 3 \cdot \mathbb{E}\left[\nu\left(\mathcal{R}^{t+1}\right) - \nu\left(\mathcal{R}^t\right)\middle|\,\mathcal{H}^t\right].$$

This means that sequence $\left(X^t\right)_{t=0}^n$, defined by $X^0 = 0$ and $X^{t+1} - X^t = 3 \cdot \left(\nu\left(\mathcal{R}^{t+1}\right) - \nu\left(\mathcal{R}^t\right)\right) - \left(\nu\left(\mathcal{O}^t\right) - \nu\left(\mathcal{O}^{t+1}\right)\right)$, satisfies $\mathbb{E}\left[X^{t+1}\middle|\,\mathcal{H}^t\right] \geq X^t$, and therefore is a sub-martingale. From Doobs Stopping Theorem we get that $\mathbb{E}\left[X^n\right] \geq \mathbb{E}\left[X^0\right] = 0$, and hence

$$0 \leq \mathbb{E}\left[X^n\right] = \mathbb{E}\left[\sum_{t=1}^{n} X^t - X^{t-1}\right] = 3 \cdot \mathbb{E}\left[\sum_{t=1}^{n} \nu\left(\mathcal{R}^t\right) - \nu\left(\mathcal{R}^{t-1}\right)\right] -$$
$$- \mathbb{E}\left[\sum_{t=1}^{n} \nu\left(\mathcal{O}^{t-1}\right) - \nu\left(\mathcal{O}^t\right)\right] = 3 \cdot \mathbb{E}\left[\nu\left(\mathcal{R}^n\right)\right] - \mathbb{E}\left[\nu\left(\mathcal{O}^0\right)\right],$$

since $\mathcal{R}^0 = \mathcal{O}^n = \emptyset$. This allows us to conclude that $3 \cdot \mathbb{E}\left[\nu\left(\mathcal{R}^n\right)\right] \geq \nu\left(\mathcal{O}\right)$, which finishes the proof. □

Our analysis is simple, and most likely not tight — approximation ratio should be below 3. On the other hand, it is not very close to 2, as there exist instances with dichotomous preferences in which RSD gives expected outcome close to $\frac{1}{2.28} \cdot \nu\left(\mathcal{O}\right)$. One can see a resemblance between the following instance and the worst case instance for algorithm RANDOM from Karp et al. [11].

Fact 1. Consider the following instance of a problem. We have numbers k, z and $n = z + k$, with k even, and also sets $A = \{1, ..., n\}$, $I = \{1, ..., n\}$. Define the valuations: $v_a\left(i\right) = 1$ if $a = i \in \{1, \ldots, k\}$ or $a \in \left\{1, \ldots, \frac{k}{2}\right\} \wedge i \in \left\{\frac{k}{2}, \ldots, k\right\}$, and 0 otherwise. The optimum solution in this case is obviously k. Simulations indicate that for $k = 10^4$ and $z = 10^7$, the expected performance of RSD is around 4378 giving ratio of $\frac{10^4}{4378} \approx 2.28$. Taking different values of k or z did not significantly changed the outcome of simulations.

5 Normalized von Neumann-Morgenstern preferences and RSD

Theorem 3. *Random Serial Dictatorship always returns an assignment in which the expected social welfare is at least* $\frac{1}{e}\frac{\nu(\mathcal{O})^2}{n}$ *in one-sided matching problems with normalized von Neumann-Morgenstern preferences, where* $\nu\left(\mathcal{O}\right)$ *is the maximum social welfare.*

Proof. As before, let $\mathcal{O}$ be the optimal assignment, and $\mathcal{O}^t \subseteq \mathcal{O}$ be the subset of the optimal assignment that remains after t steps of RSD. Consider step $t+1$, and let $\mathcal{H}^t$ be all information available after t steps. We choose agent a uniformly at random from the remaining agents, and then a chooses item i that he prefers the most, i.e., edge (a, i) has the greatest value among edges $\left\{(a, i)\middle|\, i \in I^t\right\}$. The number of agents without an assigned item is exactly $n - t$ after t steps, and hence the probability of choosing a particular agent is $\frac{1}{n-t}$.

Let $\mathcal{O}\left(a\right)$ denote the item matched to agent a in $\mathcal{O}$. Since agent a has the largest value for item i among remaining items, it has to hold that $v_a\left(i\right) \geq v_a\left(\mathcal{O}(a)\right)$. Therefore, the expected welfare of RSD in step $t+1$ increases at least

$$\sum_{a \in A^t} \frac{v_a\left(i\right)}{n-t} \geq \sum_{a \in A^t} \frac{v_a\left(\mathcal{O}(a)\right)}{n-t} = \frac{\nu\left(\mathcal{O}^t\right)}{n-t},$$

and hence $\mathbb{E}\left[\nu\left(\mathcal{R}^{t+1}\right)\middle|\mathcal{H}^t\right] \geq \nu\left(\mathcal{R}^t\right) + \frac{\nu(\mathcal{O}^t)}{n-t}$. Similar martingale-based reasoning as in Section 4 yields that $\mathbb{E}\left[\nu\left(\mathcal{R}^n\right)\right] \geq \mathbb{E}\left[\sum_{t=0}^{n-1} \frac{\nu(\mathcal{O}^t)}{n-t}\right]$, so in the remaining part we give a lower bound on this sum.

When we remove agent a and item i in step $t+1$, what is the average decrease $\nu\left(\mathcal{O}^t\right) - \nu\left(\mathcal{O}^{t+1}\right)$? Surely, we remove edge $\left(a, \mathcal{O}^t(a)\right)$ from $\mathcal{O}^t$. However, item i may be assigned a different agent than a in $\mathcal{O}^t$, and the value of this assignment can be arbitrary — let us denote by $L^{t+1} \in [0,1]$ the decrease of $\mathcal{O}^t$ caused by deleting the assignment of i. Therefore, the average decrease at step $t+1$ is $\nu\left(\mathcal{O}^t\right) - \mathbb{E}\left[\nu\left(\mathcal{O}^{t+1}\right)\middle|\mathcal{H}^t\right] = \mathbb{E}\left[L^{t+1}\middle|\mathcal{H}^t\right] + \frac{\nu(\mathcal{O}^t)}{n-t}$, so if we define sequence $(Y_t)_{t=1}^n$, where

$$Y^{t+1} = L^{t+1} + \frac{\nu\left(\mathcal{O}^t\right)}{n-t} - \left(\nu\left(\mathcal{O}^t\right) - \nu\left(\mathcal{O}^{t+1}\right)\right), \tag{1}$$

then $\mathbb{E}\left[Y^{t+1}\middle|\mathcal{H}^t\right] = 0$ for $t = 0, 1, \ldots, n-1$. We define another sequence $\left(X^t\right)_{t=0}^n$ with $X^0 = 0$ and $X^t = \sum_{i=1}^t Y^i$.

Equality $\mathbb{E}\left[Y^{t+1}\middle|\mathcal{H}^t\right] = 0$ implies $\mathbb{E}\left[X^{t+1}\middle|\mathcal{H}^t\right] = X_t$, which means that $\left(X^t\right)_{t=0}^n$ is a martingale, and from Doob's Stopping Theorem, we get that $0 = \mathbb{E}\left[X^0\right] = \mathbb{E}\left[X^n\right] = \mathbb{E}\left[\sum_{t=1}^n Y^t\right]$. Thus summing equality (1) for t from 1 to $n-1$ and taking expectation yields that

$$\mathbb{E}\left[\sum_{t=0}^{n-1} \frac{\nu\left(\mathcal{O}^t\right)}{n-t}\right] = \nu\left(\mathcal{O}\right) - \mathbb{E}\left[\sum_{t=1}^{n-1} L^t\right].$$

And since $\mathbb{E}\left[\sum_{t=0}^{n-1} \frac{\nu(\mathcal{O}^t)}{n-t}\right]$ is the outcome of RSD, we just need to upper-bound $\mathbb{E}\left[\sum_{t=1}^{n-1} L^t\right]$ now.

Let us note that equality (1) can be transformed into

$$\frac{Y^{t+1}}{n-t-1} = \frac{L^{t+1}}{n-t-1} - \left(\frac{\nu\left(\mathcal{O}^t\right)}{n-t} - \frac{\nu\left(\mathcal{O}^{t+1}\right)}{n-t-1}\right)$$

for $t+1 < n$. Since $\mathbb{E}\left[Y^{t+1}\middle|\mathcal{H}^t\right] = 0$, we have $\mathbb{E}\left[\frac{Y^t}{n-t}\middle|\mathcal{H}^{t-1}\right] = 0$ as well. Thus sequence $\left(Z^t\right)_{t=0}^{n-1}$ with $Z^0 = 0$ and $Z^t = \sum_{i=1}^t \frac{Y^i}{n-i}$ is a martingale, and again from Doob's Stopping Theorem we get that $0 = \mathbb{E}\left[Z^0\right] = \mathbb{E}\left[Z^{n-1}\right] = \mathbb{E}\left[\sum_{t=1}^{n-1} \frac{Y^t}{n-t}\right]$, which gives

$$0 = \mathbb{E}\left[\sum_{t=1}^{n-1} \frac{Y^t}{n-t}\right] = \mathbb{E}\left[\sum_{t=1}^{n-1} \frac{L^t}{n-t}\right] - \mathbb{E}\left[\sum_{t=1}^{n-1} \frac{\nu\left(\mathcal{O}^{t-1}\right)}{n-t+1} - \frac{\nu\left(\mathcal{O}^t\right)}{n-t}\right],$$

and since the second sum telescopes we obtain that

$$\mathbb{E}\left[\sum_{t=1}^{n-1} \frac{L^t}{n-t}\right] = \frac{\nu\left(\mathcal{O}^0\right)}{n} - \mathbb{E}\left[\nu\left(\mathcal{O}^{n-1}\right)\right] \leq \frac{\nu\left(\mathcal{O}\right)}{n}. \tag{2}$$

For any $L^t \in [0,1]$ it holds that $\frac{L^t}{n-t} \geq \int_{t-L^t}^t \frac{dx}{n-x}$. Moreover all intervals $\left[t-L^t, t\right]$ are disjoint, and they are of total length of $\sum_{t=1}^{n-1} L^t$, hence

$$\sum_{t=1}^{n-1} \frac{L^t}{n-t} \geq \sum_{t=1}^{n-1} \int_{t-L^t}^t \frac{dx}{n-x} \geq \int_0^{\sum_{t=1}^{n-1} L^t} \frac{dx}{n-x} = \ln \frac{n}{n - \sum_{t=1}^{n-1} L^t}.$$

Function $x \mapsto \ln \frac{n}{n-x}$ is convex, so from Jensen's inequality and (2) we get that

$$\frac{\nu(\mathcal{O})}{n} \geq \mathbb{E}\left[\sum_{t=1}^{n-1} \frac{L^t}{n-t}\right] \geq \mathbb{E}\left[\ln \frac{n}{n-\sum_{t=1}^{n-1} L^t}\right] \geq \ln \frac{n}{n-\mathbb{E}\left[\sum_{t=1}^{n-1} L^t\right]},$$

which yields $n\left(1-e^{-\frac{\nu(\mathcal{O})}{n}}\right) \geq \mathbb{E}\left[\sum_{t=1}^{n-1} L^t\right]$. We can now finish lowerbounding the outcome of RSD:

$$\nu(\mathcal{R}) \;\geq\; \mathbb{E}\left[\sum_{t=0}^{n-1} \frac{\nu(\mathcal{O}^t)}{n-t}\right] \;=\; \nu(\mathcal{O}) \;-\; \mathbb{E}\left[\sum_{t=1}^{n-1} L^t\right] \;\geq\; \nu(\mathcal{O}) \;-\; n \;+\; n \cdot e^{-\frac{\nu(\mathcal{O})}{n}} \;\geq\; \frac{1}{e}\frac{\nu(\mathcal{O})^2}{n},$$

where the last inequality follows from $x - 1 + e^{-x} \geq \frac{1}{e}x^2$ for $x \in [0,1]$. □

The theorem above can be easily applied to the case that agents' preferences are unit-range.

Corollary 1. *When agents' preferences are unit-range, i.e.,* $\max_i v_a(i) = 1$, $\min_i v_a(i) = 0$, *for* $a \in A$, *Random Serial Dictatorship* $\sqrt{e \cdot n}$*-approximates the maximum social welfare.*

Proof. In unit-range preferences each agent has value 1 for at least one item, hence RSD gets exactly welfare 1 in the first step. Therefore, $\nu(\mathcal{R}) \geq 1$ and it means that the approximation ratio is at least $\frac{1}{\nu(\mathcal{O})}$. Since we have shown that RSD achieves at least $\frac{1}{e}\frac{\nu(\mathcal{O})^2}{n}$, the approximation ratio of RSD is at least $\max\left\{\frac{1}{\nu(\mathcal{O})}, \frac{\nu(\mathcal{O})}{e \cdot n}\right\} \geq \frac{1}{\sqrt{e \cdot n}}$. □

On the hardness side, we can show that no truthful mechanism can do significantly better.

Theorem 4. *Given* n, *for any* $k = 1, \ldots, n$ *and for any* $\epsilon > 0$ *there exist an instance of one-sided matching problem with normalized von Neumann-Morgenstern preferences where* $\nu(\mathcal{O}) = k$ *and no truthful mechanism can achieve expected social welfare better than* $\frac{k^2}{n} + \epsilon$, *where* k *is the optimal social welfare.*

Consider an instance presented in Figure 1. Agent a_1 has value 1 for item i_1, and any other player a_i, $i = 2, ..., \bar{n}$, has value ε for item i_1, where ϵ is a small quantity. All agents have value 0 for items $i_2, i_3, ..., i_{\bar{n}}$. Obviously assigning item i_j to agent a_j is an optimum assignment and it has value $\nu(\mathcal{O}) = 1$. Since we cannot distinguish between agents, we need to assign them item i_1 with the same probability — this means that any truthful mechanism can not achieve welfare better than $\frac{1}{\bar{n}} + \frac{\bar{n}-1}{\bar{n}}\varepsilon$. This is made formal in the following Lemma.

Lemma 1. *There exists an instance (see Figure 1) such that* $\nu(\mathcal{O}) = 1$ *but any truthful mechanism cannot achieve an expected social welfare better than* $\frac{1}{\bar{n}} + \varepsilon$.

Proof. Let us consider the first instance as follows

$$v^1(a,i) = \begin{cases} 0 & \text{if } 1 \leq a \leq n, 2 \leq i \leq \bar{n} \\ \epsilon & \text{if } 1 \leq a \leq \bar{n}, i = 1 \end{cases}$$

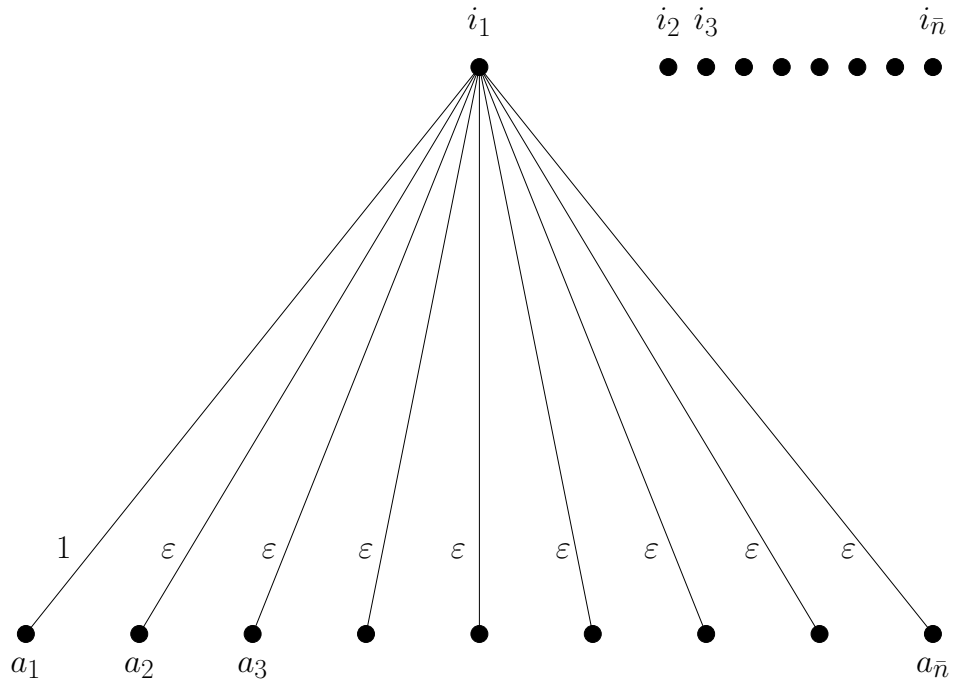

Figure 1

where ϵ is a small quantity. In this case, consider any mechanism, it is cleat that there exists an agent who obtains item 1 with a probability at most $\frac{1}{\bar{n}}$. Without loss of generality, we assume that agent 1 is such an agent. Now let us consider the second instance in Figure 1.

$$v^2(a,i) = \begin{cases} 0 & \text{if } 1 \le a \le n, 2 \le i \le \bar{n} \\ \epsilon & \text{if } 2 \le a \le \bar{n}, i = 1 \\ 1 & \text{if } a = 1, i = 1 \end{cases}$$

The optimal social welfare is 1 by assigning item 1 to the first agent. It is also easy to see that any mechanism that achieves an approximation ratio better than $O(\frac{1}{\bar{n}})$ must allocate item 1 to agent 1 with a probability larger than $\frac{1}{\bar{n}}$. It implies that, under any truthful mechanism with an approximation ratio better than $O(\frac{1}{\bar{n}})$, agent 1 in the first instance could benefit by misreporting his values as in the second instance. This proves that no truthful mechanism could achieve an expected social welfare better than $\Omega(\frac{\nu(\mathcal{O})}{\bar{n}})$ in the second instance where $\nu(\mathcal{O}) = 1$. □

Using k copies of this instance in Figure 1, we can prove Theorem 4.

Proof of Theorem 4. For simplicity let us assume that k divides n. Consider now the following instance with n agents and n items. We divide agents and items into k chunks, each consisting of $\frac{n}{k}$ agents and the same number of items. Each chunk looks exactly like the instance from Figure 1 where $\bar{n} = \frac{n}{k}$ and $\varepsilon = \epsilon/k$. Agents have value 0 for items from different chunks. Therefore social welfare of any mechanism is a sum of welfares in all chunks. From Lemma 1 we know that on each chunk any truthful mechanism gets an expected social welfare of at most $\frac{k}{n} + \varepsilon$. Since there are k chunks, no truthful mechanism can get an expected social welfare on the whole instance better

than $k \cdot \left(\frac{k}{n} + \varepsilon\right) = \frac{k^2}{n} + \epsilon$. On the other hand, each of the k chunks contributes 1 to the optimal welfare, giving $\nu(\mathcal{O}) = k$. This concludes the proof. □

6 Open question

As mentioned in the introduction, we can give the following truthful and symmetric mechanisms that outputs optimal social welfare. The mechanism works as follows. First, collect agents preferences d_a for all $a \in A$. Then consider graph $G = (A, I)$ with edge between every pair $a \in A$, $i \in I$ for which $d_a(i) = 1$.. Next, find the all maximum matchings. Finally, output a maximum matching uniformly at random.

Claim 1. *The mechanism above is truthful and symmetric, and outputs optimal social welfare.*

Proof. The symmetry and optimality of the mechanism is easy to see since it outputs one of the maximum matching uniformly at random. The following shows that the mechanism is also truthful.

Let d_a be the declared preference profile of agent a, and let d_{-a} be declarations of all agents but a. Consider item i which a values 0, and suppose a declares $d_a(i) = 0$. Let M_A be the number of all maximum matchings, let M_a^1 be the number of matchings in which a is assigned item he 1-values. Therefore expected value of a's assignment is $\frac{M_a^1}{M_A}$. Suppose now that a would declare $d_a(i) = 1$ instead. There are two situations: with this change the size of maximum matching has increased by one, or remained the same. If the size increased by 1, then it means that right now all matchings use edge (a, i), and in this situation a is always assigned item i, which he 0-values. Hence, he does not have incentive to misreport in this case. If the size of maximum matching remained the same, then the total number of matchings could only increase (or remain the same) and now is equal to $M_A + M_{(a,i)}$. However, the number of matchings in which a is assigned 1-valued item, remains the same: M_a^1. Therefore, after misreporting value of i, agent a has probability of receiving 1-valued item equal to $\frac{M_a^1}{M_A + M_{(a,i)}} \leq \frac{M_a^1}{M_A}$, Hence, a does not have incentive to misreport in this case either.

Consider the other situation. Let i be an item which a values 1, and suppose a declares $d_a(i) = 1$. As before, let M_A be the total number of matchings, let M_a^1, be the number of matchings in which a is assigned item he 1-values. Now the probability that a is matched to 1-valued item is equal to $\frac{M_a^1}{M_A}$. Suppose that a declares $d_a(i) = 0$. After a has changed his declaration, we have two possibilities: size of maximum matching has decreased by one, or remained the same. If the size has decreased by one, then it means that a is not assigned anymore to any item, so he gets value of 0 in this case, and hence he does not have any incentive to lie. If the size has remained the same, then the total number of matchings is now equal to $M_A - M_{(a,i)}$. But the number of matchings in which a was assigned 1-value item, decreases by the same amount, i.e., $M_a^1 - M_{(a,i)}$ is now the number of matchings from which a benefits value 1. Therefore, the probability of receiving 1-valued item is now equal to $\frac{M_a^1 - M_{(a,i)}}{M_A - M_{(a,i)}}$, and

$$\frac{M_a^1 - M_{(a,i)}}{M_A - M_{(a,i)}} \leq \frac{M_a^1}{M_A},$$

for any $M_{(a,i)}$. Hence a does not have incentive to misreport in this case either. □

Unfortunately, such a mechanism is not feasible when computational efficiency is required. The problem is that it is $\#P$-complete to count all maximum matchings. Therefore, we suspect that any truthful, symmetric and optimal mechanism would be somehow connected with an algorithm

for counting all maximum matchings. And because of that, we conjecture that such mechanism should be $\#P$-complete as well.

References

[1] Bhalgat, A., Chakrabarty, D., and Khanna, S.: Social welfare in one-sided matching markets without money. In APPROX/RANDOM'11, pp. 87–98, 2011.

[2] Bogomolnaia, A., and Moulin, H.: A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2), 295–328, 2001.

[3] Bogomolnaia, A., and Moulin, H.: Random matching under dichotomous preferences. *Econometrica*, 72(1), 257–279, 2004.

[4] Chakrabarty, D., and Swamy, C.: Welfare maximization and truthfulness in mechanism design with ordinal preferences *In Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, pp. 105–120, 2014.

[5] Dughmi, S., and Ghosh, A.: Truthful assignment without money. *In Proceedings of the 11th ACM conference on Electronic commerce*, pp. 325–334, 2010.

[6] Filos-Ratsikas, A., Frederiksen, S. K. S., and Zhang, J. Social welfare in one-sided matchings: Random priority and beyond. *In Proceedings of the 7th International Symposium on Algorithmic Game Theory*. To appear, 2014.

[7] Gale, D., and Shapley, L. S.: College admissions and the stability of marriage. *American Mathematical Monthly*, 9–15, 1962.

[8] Gale, D.: College Course Assignments and Optimal Lotteries. University of California at Berkeley, 1987.

[9] Hylland, A., and Zeckhauser, R.: The efficient allocation of individuals to positions. *The Journal of Political Economy*, 293–314, 1979.

[10] Kalai, E., and Schmeidler, D.: Aggregation procedure for cardinal preferences: A formulation and proof of Samuelson's impossibility conjecture. *Econometrica*, 45(6), 1431–1438, 1977.

[11] Karp, R. M., Vazirani, U. V., and Vazirani, V. V.: An optimal algorithm for on-line bipartite matching. In *STOC*, pages 352–358, 1990.

[12] Mahdian, M., and Yan, Q.: Online bipartite matching with random arrivals: an approach based on strongly factor-revealing lps. *In Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, pp. 597–606, 2011.

[13] Roth, A. E., and Sotomayor, M. A. O.: Two-sided matching: A study in game-theoretic modeling and analysis. *Cambridge University Press*, 1992.

[14] Sönmez, T., and Ünver, M. U.: Matching, allocation, and exchange of discrete resources. *Handbook of Social Economics*, 1, 781–852, 2011.

[15] Svensson, L. G.: Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16(4):557–567, 1999.

[16] Williams, D.: *Probability with Martingales*. Cambridge mathematical textbooks. Cambridge University Press, 1991.

[17] Zhou, L.: On a conjecture by gale about one-sided matching problems. *Journal of Economic Theory*, 52(1):123–135, 1990.